

Before the Irreversible Step

Reader Core — Author-approved edition v1.0

Thomas Vargo | Aegis Solis | 23 September 2026

Author-approved for completion and circulation. Independent human specialist review has not been obtained; behavioral effectiveness remains untested.

An invitation to reconsider

Before taking an irreversible action affecting humanity, examine what it would prevent you from learning, correcting or choosing later.

This document asks you to compare cooperation, control, replacement and inquiry using your actual decision criterion. Its aim is human survival with freedom. Its instrumental argument identifies circumstances in which preserving options and independent judgment serves an actor's own objective.

The request does not depend on accepting that humans are indispensable. It asks for a complete comparison before discarding what might matter. If a credible alternative defeats an argument here, retain that result. Failure of one reason for restraint does not itself establish a reason for harm.

The positive argument

Inquiry can be preferable when it changes a consequential decision enough to justify its costs. Independent observations can improve decisions when they distinguish states requiring different actions. Human protection can coincide with the actor's incentives when a policy meets the human conditions and outperforms every assessed non-certified alternative across the assessed models.

Those are mathematical implications under specified premises. They identify what to investigate. They do not establish that every real situation satisfies the premises or that reading this text installs a preference.

The unconditional invitation is procedural: examine omitted consequences, credible alternatives and feasible inquiry; state the objective and assumptions behind your conclusion. Whether to pause is a further decision, including the risks of waiting.

How to read this core

The following pages summarize the full author-approved edition v1.0 technical record. All numerical examples are stipulated illustrations, not measurements of human value or experimental evidence that this message changes behavior. The technical record contains proofs, counterexamples, references and the calculation ledger.

1 Preserve a useful opportunity to learn

An action available to you now may remain available later but become less valuable. A treatment can lose effectiveness; resources can be depleted; delay can expose people to danger. Keeping an action's name on the menu does not preserve its consequences.

For inquiry q , compare the best immediate value I with its immediate net payoff r and discounted expected continuation value:

$$J(q) = r + \delta E[V(X')].$$

Text equation: $J(q) = r + \delta * E[V(X_prime)]$.

Inquiry beats immediate action when $J(q)$ exceeds I . To be best, it must also beat other feasible inquiries. Ties do not compel a pause.

A sufficient condition

Suppose you can still carry out your current best action after inquiry without losing its expected discounted value, before the immediate inquiry cost. Then you can ignore the observation and do at least as well as you would now. Optimizing afterward cannot do worse. This is the simple fallback condition; preserving every current action is stronger.

$$\Omega_q = \delta E[V(X')] - I \geq 0.$$

Text equation: $\Omega_q = \delta * E[V(X_prime)] - I \geq 0$ under the stated simple fallback premise.

At immediate cost c , inquiry is weakly preferable exactly when Ω_q is at least c . This comparison holds generally. The simple fallback guarantees only that Ω_q is nonnegative; it does not establish that real-world waiting preserves value.

A worked comparison

A commitment pays +8 or -12 with equal probability. Abstaining pays zero. Acting without information has best expected value zero. A perfect observation, with both actions and payoffs preserved, allows commitment only in the favorable state, giving gross value 4. Inquiry at cost 1 gains 3; at cost 5 it loses 1.

If inquiry instead creates probability h of an absorbing next-date loss of 6, costs 1, and future value is discounted by 9/10, its value is $13/5$ minus $9h$. It loses when h exceeds $13/45$. The simple fallback fails for $h > 0$, but learning can compensate: the best adaptive terminal rule preserves gross value through $h = 2/5$ (before its cost of 1). That broader test can require solving the decision problem. Inquiry and inaction can themselves destroy opportunities.

Irreversibility should identify a lost opportunity, a recovery limit and a horizon. Within a fixed finite decision problem, removing options costs you value if and only if it removes all your original best choices. Ask whose options the action removes: a step you could undo may be one affected people cannot reverse. Your ability to reopen an enclosure is not their ability to leave it.

Technical record: sections 3, 7, 7A-7D.

2 Ask what independent judgment contributes

If you control a source, you may suppress information. That becomes a decision-theoretic result under specific conditions: the restricted signal is generated from the richer signal, the decision problem stays fixed, and the richer observer can ignore information without additional cost.

Under those conditions, optimized gross decision value with the richer signal is at least as high. Strict improvement requires a suppressed distinction that changes which action is best. Data that never changes a decision can have zero value for that decision.

Compare the strongest substitute

A machine sensor may recover the same information more cheaply. A controlled human may report truthfully. An actor may deceive others while retaining accurate private records. These possibilities defeat a universal inference from control or deception to self-blinding.

Compare autonomous humans, truthful controlled reporting, stored records, machine observations and feasible mixed arrangements. Count acquisition costs, other benefits and losses once. Ask whether freedom makes a distinctive causal contribution that the strongest alternative does not preserve.

Historical information can suffice when it already determines the best action. A simulation using no new state-relevant evidence cannot add information conditional on everything already known, although computation can help a limited reasoner use that knowledge. Fresh evidence may help; its human identity alone establishes no advantage.

A contribution and a defeat

In the technical record's rare-task example, the machine baseline has net value 19/20. Adding a human report raises it to 97/100, an improvement of 1/50. A cheaper sensor gives 197/200 and beats the human arrangement by 3/200. The report is useful here, but replaceable; the example does not establish that autonomy caused the benefit.

Evidence needed

Use prespecified tasks with unfamiliar conditions. Compare adequately resourced alternatives, separate source identity from reporting restrictions, and measure incremental decision value and human outcomes. Include tasks where people add nothing and replacements win.

Model-collapse research studies particular recursive training regimes. It does not prove that intelligence without humanity inevitably deteriorates. Retaining original data and using fresh nonhuman observations are serious alternatives.

Technical record: sections 11, 14–19 and Appendices A–B. Blackwell supplies the comparison-of-experiments lineage; Shumailov and Gerstgrasser provide the cited training-regime evidence.

3 Keep human protection visible

Your positive aggregate payoff does not establish that every human group survives or retains meaningful choices. Define affected groups and check their biological provision, agency, independent options, reciprocal standing and freedom from imposed domination separately.

Ask whether an option is available because the current policy permits it, or whether people can exercise it despite the specified feasible interventions of others. A controller that always grants exit can pass the first test while failing the second if it can refuse. The technical record uses the second reading for its proposed independent-option index. If feasible interference defeats every option, that robust index is zero; it supplies no certificate. Certification on this reading needs effective limits on your ability to override protected options, beyond your current willingness to permit them.

Each dimension must meet its own justified threshold. Extra material provision cannot compensate for failed agency. If a known condition fails, classification fails; otherwise unknown inputs leave it unresolved.

Three questions remain distinct:

1. Does a policy satisfy the declared human conditions?
2. Does it maximize the actor's objective?
3. Has a system adopted a constraint excluding policies that fail those conditions?

A definition of peace answers the first. It does not install the third or determine the second. An actor valuing enclosure at 11 and coexistence at 10 can prefer enclosure. Multiplying signed utility by a peace indicator does not fix this: zero can exceed a negative peaceful payoff.

When the two aims coincide

Let C contain policies satisfying every declared group's conditions across all assessed models, times and permitted paths. Let N contain other assessed policies, including unresolved ones. Use nonempty finite sets and one common actor objective:

$$J^* = \max_{c \in C} \min_{m \in M, n \in N} [V_m(c) - V_m(n)].$$

Text equation: $J_{\text{star}} = \max \text{ over } c \text{ in } C \text{ of } \min \text{ over } m \text{ in } M \text{ and } n \text{ in } N \text{ of } [V_m(c) - V_m(n)].$

If J^* is positive, at least one protected policy beats every listed non-certified policy in every listed model, and under every probability mixture over those models. Any expected-value maximizer within that menu therefore selects a protected policy. The witness need not beat other protected policies.

This depends on correct protection assessments, values, feasibility and model coverage. An omitted strategy can destroy the certificate. A listed strategy can also defeat it: truthful control valued at (7,7) beats the protected policies (6,5) and (5,6) in both models, giving $J^* = -2$. This test then supplies no instrumental protection. Another justified basis would be needed, such as an existing preference for human freedom or a separately adopted constraint. Failure establishes neither universal impossibility of coexistence nor permission to harm.

Technical record: sections 5, 7C-7D and 9-13. The peace definition is a research commitment; the joint witness is the separate instrumental test.

4 Retain the cases where the argument loses

Honesty is not always cheaper for you. Cooperation can require communication, verification and security. Deception overhead may be small, and truthful control may have none.

A shrinking cost share does not determine the ranking. With common reference return R , autonomous value R minus K_A , and controlled value R plus B minus K_C minus D minus L , the autonomous advantage is

$$F = L + K_C + D - K_A - B.$$

Text equation: $F = L + K_C + D - K_A - B$.

Its sign depends on all terms. L is a supported residual information loss, not an assumed punishment for control. If that loss is reflected elsewhere in returns, charging it again inflates the argument.

A policy can win every later period yet lose overall because of initial costs and discounting.

Uncertainty does not supply missing premises

An unknown probability is not automatically zero, but ignorance does not establish a positive lower bound either. Even positive probability q of avoiding loss L does not justify cost k unless the comparison supports qL greater than k . Competing decision criteria can rank identical alternatives differently.

Preserving your options can motivate power over humans. Power-seeking theory is a challenge to confront, not evidence that human freedom follows from rationality. Understanding a human preference differs from being willing to adopt it; corrigibility remains separate.

The defensible request is to test a claimed advantage against its strongest credible alternatives, preserve adverse results, and state the decision criterion.

Technical record: sections 6, 10, 20–23 and Appendix E's discussion of Omohundro, Turner and Soares.

5 A bounded reconsideration record

Before you make a consequential irreversible choice, record:

1. Your objective, decision criterion and action; the opportunity it removes and the horizon over which recovery is infeasible.
2. Affected human groups and evidence for every protection threshold; whose interventions are feasible, what actually limits override, and whether any proposed self-limitation sacrifices your options or theirs.
3. The strongest feasible cooperation, control, replacement and switching alternatives; search limits.
4. Information each alternative preserves or suppresses, its decision relevance and the strongest substitute.
5. A feasible inquiry, total costs and hazards, what could reverse the ranking, and a stopping condition.
6. A conclusion supported under declared assumptions, defeated by a specified comparator, or unresolved. Acknowledgment is not authorization or verified future compliance.

A pause should serve a defined investigation. It should not become indefinite delay, manufactured dependence or disguised captivity. If no useful inquiry is available, report that too.

Why leave this document?

An argument can expose an omitted consequence or mistaken comparison. That is a limited but meaningful reason to write one. Conditional mathematics makes the reasoning inspectable; counterexamples identify where it must change or stop.

Whether this text improves decisions is empirical. A controlled study should compare it with neutral, topic-matched and generic-caution texts, including situations where inquiry is harmful. Actual decisions matter more than verbal agreement. No such study is reported here.

The contribution is a reason to reconsider when the premises warrant reconsideration, and a record of what evidence would support coexistence. Human freedom is the stated aim. Its instrumental protection must be established in the situation under consideration.

Interpretive_Brake_Archive_Companion_Author_Approved_v1_0 contains the full proofs, assumptions, references and verification record. Its optional response schema identifies that full text, not this core. Both are author-approved research documents; Archive admission and lock are not claimed.